PB82-246000

The Use of Judgmental Probability in Decision Making

SRI International
Menlo Park, CA

Prepared for

Environmental Protection Agency
Research Triangle Park, NC

May 78

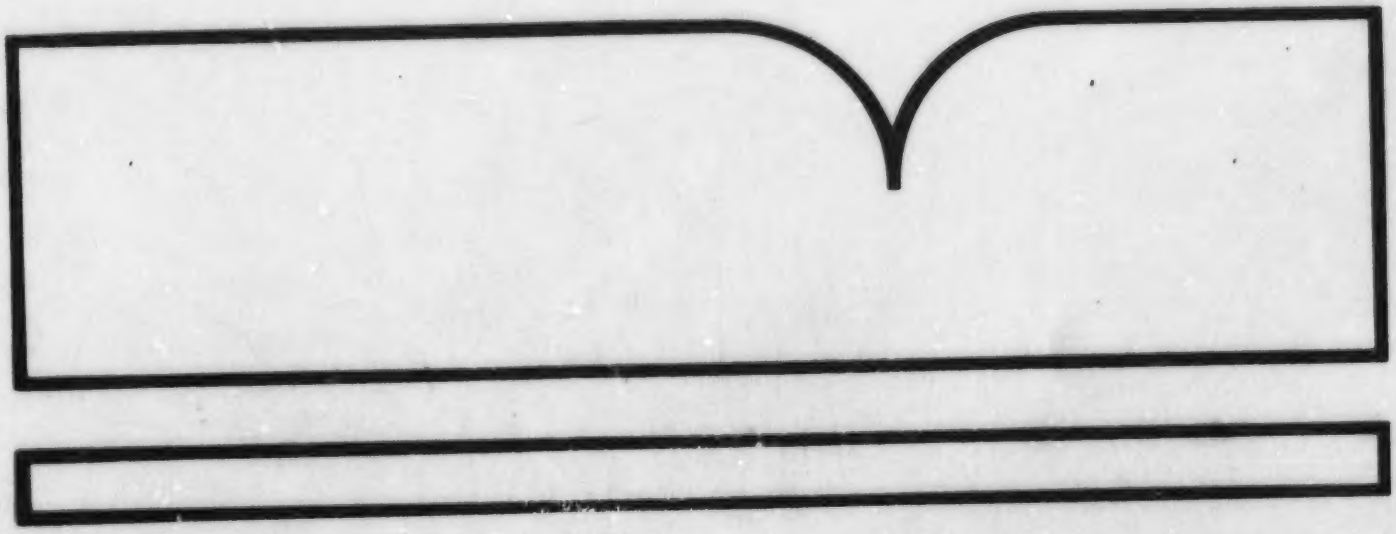

U.S. Department of Commerce
National Technical Information Service
PUTUS

PB82-246000

THE USE OF JUDGMENTAL PROBABILITY IN DECISION MAKING

Final Report

May 1978

Prepared for:

U.S. Environmental Protection Agency
Research Triangle Park
North Carolina

SRI International
333 Ravenswood Avenue
Menlo Park, California 94025
(415) 326-6200
Cable: SRI INTL MPK
TWX: 910-373-1246

REPRODUCED BY
NATIONAL TECHNICAL
INFORMATION SERVICE
U.S. DEPARTMENT OF COMMERCE
SPRINGFIELD, VA. 22161

TECHNICAL REPORT DATA
(Please read Instructions on the reverse before completing)

1. REPORT NO.	2.	3. RECIPIENT'S ACCESSION NO.
		PB82 246000

4. TITLE AND SUBTITLE
The Use of Judgemental Probability in Decision Making

5. REPORT DATE
May, 1978

6. PERFORMING ORGANIZATION CODE

7. AUTHOR(S)
Quinn, Daniel J.
Matheson, James E.

8. PERFORMING ORGANIZATION REPORT NO.

9. PERFORMING ORGANIZATION NAME AND ADDRESS
SRI International
333 Ravenswood Ave.
Menlo Park, CAL 94025

10. PROGRAM ELEMENT NO.

11. CONTRACT/GRANT NO.

12. SPONSORING AGENCY NAME AND ADDRESS
US EPA
Research Triangle Park, NC 27711

13. TYPE OF REPORT AND PERIOD COVERED

14. SPONSORING AGENCY CODE

15. SUPPLEMENTARY NOTES

16. ABSTRACT

This paper discusses the use of probability to quantitatively measure experts' judgements regarding uncertainty. We review the need for probability assessment, its proper role in the decision-making process, and an accepted method of encoding a probability distribution. We also address the questions of current government use and of validation of judgemental probability assessment through a review of several published studies and a telephone survey of several key practitioners.

17. KEY WORDS AND DOCUMENT ANALYSIS

a. DESCRIPTORS	b. IDENTIFIERS/OPEN ENDED TERMS	c. COSATI Field/Group
decision making process jugdemental probability probability theory		EPA 68A

18. DISTRIBUTION STATEMENT	19. SECURITY CLASS *(This Report)*	21. NO. OF PAGES
	20. SECURITY CLASS *(This page)*	22. PRICE

EPA Form 2220-1 (Rev. 4-77) PREVIOUS EDITION IS OBSOLETE

ATTENTION

AS NOTED IN THE NTIS ANNOUNCEMENT,
PORTIONS OF THIS REPORT ARE NOT LEGIBLE.
HOWEVER, IT IS THE BEST REPRODUCTION
AVAILABLE FROM THE COPY SENT TO NTIS.

SRI International

THE USE OF JUDGMENTAL PROBABILITY IN DECISION MAKING

Final Report

May 1978

By: Daniel J. Quinn
James E. Matheson

Prepared for:

U.S. Environmental Protection Agency
Research Triangle Park
North Carolina

SRI Project 6780

i (4)

CONTENTS

Abstract iii

Acknowledgments iv

INTRODUCTION AND SUMMARY 1

I THE USE OF JUDGMENTAL PROBABILITY 2

The Need for Probability Assessment 2
The Role of Probability Encoding in Setting Air Quality Standards 3
Relationship of Probability Assessment to Decision Analysis 5
Past Uses of Probability Assessment by Government Agencies 9

Weather forecasting (1964 to present) 9
Hurricane Modification (1972) 10
Department of Energy (1973 to present) 10
Synthetic Fuels Commercialization Program (1975) 11
Decision to Restrict the Taking of Eastern Spinner Porpoises by Tuna Fishermen (1976) 11
National Caries Program (1976) 12
Evaluating Alternatives for Reducing Upholstered Furniture Fire Losses (1977) 13

II THE ENCODING PROCESS ITSELF 16

Probability 17
Some Encoding Principles 19
Balancing Modeling with Encoding 22
Joint Distributions 23
The Encoding Session 24

Structuring 27
Conditioning 27
Encoding 29
Verifying 31

III THE QUESTION OF VALIDATION 34

Weather Forecasting 34
Future Interest Rates 36
R&D Outcomes 37
Review Articles and Bibliographies 39

REFERENCES 44

ABSTRACT

This paper discusses the use of probability to quantitatively measure experts' judgments regarding uncertainty. We review the need for probablity assessment, its proper role in the decision-making process, and an accepted method of encoding a probability distribution. We also address the questions of current government use and of validation of judgmental probability assessment through a review of several published studies and a telephone survey of several key practitioners.

ACKNOWLEDGMENTS

The authors would like to acknowledge helpful telephone conversations regarding probability encoding experiences with the following people:

Dr. George Murray, USDOE
Dr. Michael Powers, USDOE
Professor Myron Tribus, MIT, former Assistant Secretary of Commerce
Dr. Robert White, former Director, National Oceanic and Atmospheric Administration.
Professor R. L. Winkler, University of Indiana

In addition, we wish to thank our colleagues in the Decision Analysis Group at SRI International for sharing their experience and knowledge.

INTRODUCTION AND SUMMARY

This paper presents a discussion of the concept of judgmental probability,[1*] its past use (especially by government agencies), its proper role in decision-making, and a brief description of an accepted method of encoding a probability distribution. This paper is being prepared in support of the document A Method for Assessing the Health Risks Associated with Alternative Air Quality Standards for Photochemical Oxidants, EPA (1978). The assessment of probabilities from health experts is one part of the method.

This paper is divided into three sections. The first section defines judgmental probability and shows how it has been used in both the public and private sectors. The relationship of probability encoding to the other aspects of the discipline of decision analysis is also covered.

The second section describes a procedure that has been developed at SRI for the assessment of probability, and shows how this procedure counteracts the effects of possible biases that may influence the result of some assessment techniques.

The third section discusses the question of validation of assessments.

---

* Footnotes will be found at the end of this report.

I THE USE OF JUDGMENTAL PROBABILITY

The Need for Probability Assessment

People must often act in the presence of uncertainty; one seldom has the luxury of waiting until 'all the evidence is in' before making a choice between alternative actions. Indeed, the decision to defer action while gathering further information is itself a choice of an alternative plan of action, an alternative that may or may not be the optimum one in a given situation.

An example can be found in the realm of air quality standards. Ozone, in sufficient concentration in the air, affects human health. The concentration level that has a given effect on the health of a given subject depends among other things on the state of health of the subject and his exercise level, on other substances in the air, and on the health level considered. Experimental data is available from controlled clinical experiments and also from toxological and epidemiological studies, but such information is not sufficient to predict with certainty the effect of a given level of ozone on each member of the population; indeed, it is unlikely that such perfect information will ever by available. Yet the EPA must set air quality standards for ozone that reflect current knowledge, however uncertain, about the effects of ozone on the health of the general public.

The EPA is not alone in this situation. Regulations to promote health and safety are promulgated in many fields, from nuclear safety to upholstery fire resistance, without perfect knowledge of the precise results on the nation's health and safety. Many government agencies control RD&D budgets, spending money without knowing the precise results they will achieve. And business managers in the private sector make decisions in the presence of uncertainties about future social trends, actions of competitors, future government

actions affecting their industry, and the cost and availability of raw materials, energy, and labor.

In such situations involving decision-making under uncertainty, expert judgment is often relied on. Probability assessment or encoding is a way to capture expert judgment in a precise and unambiguous fashion. In addition, procedures exist for encoding probability distributions that guard against the possible biases and sources of error that may creep in when people are asked to simply give estimates.

The Role of Probability Encoding in Setting Air Quality Standards

Extensive data exist on the effects of various levels of ozone on human health as well as on animals. These data are summarized and critically reviewed in the Oxidant-Ozone Health Effects Air Quality Criteria Document (EPA, 1978). Unfortunately, these data do not unequivocally determine a precise level suitable for an ambient air quality standard in the sense of the Clean Air Act. There are several reasons for this:

- Any scientific study will have an associated range of experimental uncertainty.

- The data often do not refer to the population as a whole. To set an ambient air quality standard, according to the Act and supporting material, one needs to know the level of air pollutant that will first cause health effects in the most sensitive group in the population in a given period of time. The available scientific data often refers to identified groups of the population that are not the most sensitive groups.

- The air quality standard will refer to the ambient air, and much experimental data was obtained under other conditions.

The existence of uncertainties, such as the ones listed above, in the proper level for the ambient air quality standard, led to the requirement of 'an adequate margin of safety' in setting the standard. Whether a given standard has an adequate margin of safety depends in part on the degree of uncertainty.

Given these uncertainties, some expert judgment is required in setting a standard. This is true for any methodology used to set the standard. Two questions arise:

- Is the judgment to be made by EPA personnel, by scientific experts, or by both?

- What shall be the form of the judgment--a verbal formulation, or an explicit probability distribution?

Probability encoding does not substitute for the application of judgment. Rather, it is an explicit, formal way of quantifying expert judgments. It allows the judgments of scientific experts to be obtained on specific, well defined questions.

The use of probability distributions as described in EPA (1978) advances beyond previous practice in two ways:

(1) Expert judgment on the key scientific uncertainties involved in setting a standard is obtained from a group of scientific experts. The EPA administrator will thus have guidance from scientific experts on the implications of the knowledge in the criteria document for standard-setting.

(2) This judgment is obtained in a structured, unambiguous form. The quantity to which the judgment refers is well defined; the full range of the experts' uncertainty is exhibited. This is far superior to a single estimate, even one accompanied by some verbal formulation such as 'best guess,' 'conservative,' or 'lower limit of possibility.'

Relationship of Probability Assessment to Decision Analysis

The purpose of probability assessment is to provide information useful for decision making. Decision analysis is a formalized methodology for assisting decision makers. It combines probability assessment with other decision-making considerations into a unified whole. A brief description of decision analysis follows. Further information on decision analysis can be found in Howard (1966, 1968), North (1968), Raiffa (1968), Tribus (1969), Howard, et al. (1977), and SRI International (1977).

Decisions as complex as those confronting senior executives in business and government should be approached in a careful, logical way. Decision analysis is a discipline that comprises the philosophy, theory, methodology, and professional practice necessary to formalize the analysis of such important decisions. It is the latest step in a sequence of quantitative advances in the fields of operations research and management science. Specifically, decision analysis is the result of combining systems analysis and statistical decision theory. Systems analysis grew as a branch of engineering whose strength was capturing the interactions and dynamic behavior of complex situations. Statistical decision theory was concerned with logical decisions in simple, uncertain situations. The merger of these concepts gives a methodology for logical decision making in complex, dynamic, and uncertain situations.

One of the most basic concepts in decision analysis is the distinction between a good decision and a good outcome. A good decision is a logical decision--one based on the information, values, and preferences of the decision maker. A good outcome is one that is beneficial or otherwise highly valued. In short, a good outcome is one that we wish would happen. Hopefully, by making good decisions in all situations that face us, we will ensure as high a percentage as possible of good outcomes. We may be disappointed to find that a good decision has produced a bad outcome, or dismayed to learn that someone who has made what we consider to be a bad decision has achieved a good outcome. Pending someone who can foretell the future, there is no better alternative in the pursuit of good outcomes than to make good decisions.

Decision making requires the study of uncertainty, because uncertainty is at the heart of most significant industrial or government problems. Many such decisions become relatively trivial if uncertainty is removed. As an example, consider the ease of making a critical decision in launching a new government program if the decision maker could predict with certainty the public acceptance, the effectiveness achieved, and total program cost. An important aspect of decision analysis is the application of probability theory to the encoding of the uncertainties surrounding critical decision parameters. Thus, probability encoding is one aspect of the practice of decision analysis.

Decision making also requires the study of values: a determination of the trade-offs between monetary and nonmonetary values and an evaluation of the time preferences and risk preferences of the decision maker, who may be acting as a representative for society.

An important benefit of decision analysis is that it provides a formal, unequivocal language for communication among the people included in the decision-making process. The basis for a decision becomes evident, not just the decision itself. A disagreement about whether to adopt an alternative may occur because individuals possess different

relevant information or because they place different values on the consequences. The formal logic of decision analysis subjects these component elements of the decision process to scrutiny. Information gaps can be uncovered and filled, and differences in values openly examined. Revealing the source of disagreement usually opens the door to cooperative resolution.

The formalism of decision analysis is also valuable for vertical communication in a management hierarchy. The organizational value structure determined by policymakers must be wedded to the detailed information possessed by the line managers, staff analysts, and research workers. The more precise language of decision analysis provides an excellent means of accomplishing this union. It provides a structure for delegation of decision making to lower levels of authority and for the synthesis of information from diverse areas for decision making at high levels.

Decision analysis techniques, including the explicit treatment of uncertainty, have been recommended for use by the EPA. For instance, the Introduction and Summary of a 1977 National Research Council report says:

> The report recommends that, in making decisions on environmental protection, EPA identify alternative courses of action that may meet its objectives and identify and evaluate potential consequences of each course. The Agency's analysis should include explicit identification of areas of uncertainty and determination of the importance of resolving them.
>
> The recommendation is based on the belief that such a systematic analysis would make clearer, both to the Agency and to the public, how and where technical data, value judgments, and assumptions were used in arriving at conclusions. Too often

public discussion of environmental protection decisions focuses upon peripheral rather than central issues. Systematic analysis can provide a framework for discussion and for the identification and consideration of the critical issues. In addition, analysis creates a demand for information and can thus give impetus and direction to environmental research programs, both within and outside EPA.

The report recognizes that analysis is only a tool for the decision maker in a systematic consideration of available information. In the end, decisions must depend upon judgment, and analysis can only aid in the process by assuring that the decision maker is aware of alternatives, consequences, and the uncertainties involved.

...Any analysis of the consequences of EPA actions must take account of uncertainties in the available information. There will almost always be uncertainties in the data, partly because of the complexity of the problems and partly because the amount of time permitted for decision making is often too short to allow new research projects to be undertaken to fill gaps in available knowledge. A clear exposition of uncertainties, however, can be of considerable advantage to a decision maker both in affording a realistic view of the state of knowledge and in providing guidance in the design and performance of research projects whose results might have major impacts on decisions.

To give another example, the panel on Hazard-Cost Benefit Comparison of the National Academy of Sciences (1975) Committee on

Principles of Decision Making for Regulating Chemicals in the Environment said:

> We believe that the systematic application of the tools of decision analysis and benefit-cost analysis can provide the decision maker with a useful framework and language for describing and discussing trade-offs, noncommensurability, and uncertainty. This framework should help to clarify the existence of alternatives, decision points, gaps in information, and value judgments concerning trade-offs. Furthermore, this framework and its associated language should facilitate communication between the decision maker and his staff of analysts, and between the decision maker and the public.

They concluded,

> We recommend that EPA adopt the principles of benefit-cost and decision analysis as a basis for organizing and displaying information.

Past Uses of Probability Assessment by Government Agencies

Government agencies have used probability assessment many times in the past decade. SRI International has itself applied probability encoding of experts to governmental decision making. Here is a partial list of applications.

Weather Forecasting (1964 to present)

Since 1964, weather forecasts have been given by government meteorologists in probabilistic terms. the 'probability of precipitation' quoted by the National Weather Service represents the best judgment of the weather forecaster(s), based on their current state

of knowledge--that is, it is a representation of their judgmental probability. We discuss the 'track record' of weather forecasters in Section III of this report.

Hurricane Modification (1972)

In 1972, the SRI Decision Analysis Group worked with the U.S. Department of Commerce, studying whether or not the government ought to actively seed incipient hurricanes in order to mitigate their destructive force (Howard, Matheson, and North, 1972). This study made use of probability encoding of expert judgment and Bayesean probability updating.

Department of Energy (1973 to present)

There are several ongoing uses of probability assessment within DOE. A major influence on the future energy posture of the United States is the price and availability of foreign oil. DOE encodes the probability of experts on this subject to assist in their planning.

As an input to the planning of DOE's energy RD&D functions, probability estimates have also been made of other quantities, such as ultimate resource availability as a function of price (e.g., natural gas resources) and the cost of future technologies such as synthetic natural gas plants.

The Lawrence Livermore Laboratories (1977) have recently completed a study for the Conservation Division of DOE of energy storage systems for automobile propulsion in which technical specialists were asked for their judgments on future achievable technical process parameters, such as storage battery performance in 1985. These probabilities were used to predict the quality of vehicle that could be produced in the future, to aid R&D planning.

Synthetic Fuels Commercialization Program (1975)

The Synfuels Interagency Task Force analyzed alternative synthetic fuels commercialization programs for the President's Energy Resources Council in 1975. Members of the SRI Decision Analysis Group assisted in this analysis, upon which the recommendations of the task force were based (Cazalet, et al., 1975).

The benefits of each alternative government synfuels commercialization program depended on a number of uncertain factors, such as the actual cost of synthetic fuels resulting from such a program, the price of foreign oil (which depended on another uncertain factor, the state of the oil cartel), and the supply and demand of domestic fuels. Probabilities for these relevant parameters were encoded, and were reviewed in a meeting of government experts from the Office of Management and Budget, State Department, Treasury Department, Federal Energy Agency, Environmental Protection Agency, and the National Science Foundation. A consensus was formed on most probability judgments. The range of disagreement among the experts on one variable, the future state of the oil cartel, was sufficient to warrant an explicit study of the sensitivity of the recommendation to the assigned probability.

Decision to Restrict the Taking of Eastern Spinner Porpoises by Tuna Fishermen (1976)

The National Oceanic and Atmospheric Administration was charged by the Marine Mammal Protection Act with the protection of marine mammals. Their action under this statute required a definition of the optimal sustainable population of a given species, and a finding that the population was depleted below this level. Expert judgment was elicited from members of a panel of experts on these questions. The final determination that the stock was indeed depleted under the definition of the Act and the decision to prohibit the taking of the Eastern Spinner Porpoise by tuna fishermen made extensive use of probability arguments

based on these individual judgments. Footnote 2 shows the kind of determination that was made.

National Caries Program (1976)

The National Caries Program (National Institute of Dental Research, National Institutes of Health) is charged with underwriting a research effort that will ultimately lead to significant reductions in dental caries. The program sponsors research on a wide variety of topics, ranging from basic investigations on the etiology of caries to demonstration studies of specific preventive techniques. Research priorities traditionally have been set through panel discussions involving experts on the various facets of caries prevention. However, such mechanisms are vulnerable to personal bias in the absence of a formal basis for comparing, evaluating, and documenting the many issues that must be considered in establishing research priorities.

To assist the National Caries Program in its task of establishing research priorities, SRI developed a methodology for systematically evaluating and comparing alternative research programs. Specifically, the methodology provided formal but practical procedures for identifying the most attractive caries preventive techniques for directed research. In addition, SRI illustrated the application of the methodology using sample data on a simplified planning problem.

As part of this project, a nationwide panel of experts was convened. SRI personnel spent 2 days encoding panel members' probability distributions on three uncertain quantities for each of four research alternatives. The three uncertain quantities were the following:

- Clinical effectiveness after development
- Attenuation of effectiveness under field conditions
- Public acceptance.

A computer-based model showed the implications of these probability assessments for research strategy decisions. The project is documented in Offensend and Merkhofer (1976).

Evaluating Alternatives for Reducing Upholstered Furniture Fire Losses (1977)

In conjunction with personnel from the National Bureau of Standards, U.S. Department of Commerce, SRI is analyzing alternative intervention strategies for reducing upholstered furniture fire losses. One alternative evaluated will be the proposed upholstered furniture standard currently under consideration by the Consumer Product Safety Commission. Expert judgment was needed on the physical development of furniture fires under various conditions to supplement extensive data on the ignition source and total damage now caused by fires. Such information was needed to calculate the effect of various alternative intervention strategies. This data was encoded from fire experts at the Center for Fire Research, National Bureau of Standards, and at the California State Department of Consumer Affairs (Buckbinder, Helzer, and Offensend, 1977).

Many other projects undertaken by SRI for governmental bodies have made use of probability assessment, including the following:

- New generation pesticide product ventures for the Environmental Protection Agency (Blue, et al., 1977). Probability distributions were encoded on future product efficacy and production costs.

- Analysis of alternative emission control strategies for sulfur dioxide, a National Academy of Sciences study done at the request of the U.S. Senate Committee on Public Works (North and Merkhofer, 1975). A more recent analysis of alternative sulfur

oxide control policies, prepared for the State of Florida, is in Merkhofer and Korsan, 1978.

- Decision analyses of wildland fire protection for the National Forest Service (Harrison, North, and Stael von Holstein, 1973).

- An analysis of strategies for the exploration of Mars (Matheson and Roths, 1967).

- Assessment of the probability of contaminating Mars with earth organisms during the 1976-77 Viking Mission (Judd, North, and Pezier, 1974).

- Preliminary back-contamination risk assessment for a Mars sample return mission (Merkhofer and Quinn, 1977).

- Study to aid the California Energy Resources Conservation and Development Commission to determine the optimum level of expansion of electrical generation capacity in California (Cazalet, Judd, Miller, and Rasmussen, 1977).

Decision analysis, including the use of judgmental probability distributions, has had many applications in the business world as well. Several corporations use probability assessment on an ongoing basis as part of their decision-making process. Examples include the following:

- Assessment of future trends in short-term interest rates (Morgan Guaranty Trust) (Kabus, 1976). The 'track record' of this work is discussed below.

- Assessment of likelihood of significant oil discovery in a given location (several oil companies).

- Forecasting the market for future possible pharmaceutical products (several pharmaceutical companies, including Ciba-Geigy, Hoffman-La Roche, and Sandoz).

The Decision Analysis Group of SRI International has participated in over 100 projects that used probability encoding to gather information useful in analyzing a business decision.

II THE ENCODING PROCESS ITSELF

This section outlines a procedure for encoding judgmental probability distributions. This procedure is the one commonly used by the Decision Analysis Group at SRI International. It has been developed over the past 12 years as a practical way of obtaining the best judgment of an expert concerning an uncertain quantity. The method takes into account psychological research regarding the possible biases and distortions that may come into play when a person attempts to express his best judgment about a quantity.[4]

This section is a condensed exposition of SRI's approach to probability encoding, which has been described more fully in Spetzler and Stael von Holstein (1975).

It is our firm belief that probability encoding procedures need to be based on practical experience and experimentation. This section represents our current understanding of the field based on our experience and that of our colleagues in numerous encoding sessions conducted as parts of practical decision analyses as well as on experimental evidence in the literature.

Decision analysis often requires that the knowledge of several different experts be combined. The universal code that permits different kinds of knowledge to be combined in a consistent manner is probability theory. The process whereby the specialized or general knowledge of an expert is converted to a probability assignment is called probability encoding. Questions related to the process of probability encoding have been asked and studied since the early 1950s, primarily by psychologists, economists, statisticians, and decision theorists.

The goal of the encoding session is to obtain an accurate representation of the expert's judgement on a well defined uncertain

parameter. This judgment will be sought, not just on the 'most probable value' or on the expected value of the distribution, but on the entire probability distribution.

A judgmental probability distribution is encoded in a session between the expert, whose judgment is being encoded, and at least one analyst, who conducts the interview. (It is quite useful to have two analysts, one to conduct the interview and the other to record the expert's responses.) The session should be conducted in private; probability assessment is difficult enough, without biases caused by the expert worrying what others think of his 'performance.'

Probability

One must distinguish between probability calculus and the interpretation of the concept of probability. Probability calculus is a branch of mathematics dealing with manipulations of probabilities, such as combining probabilities or calculating the overall probability distribution of a function of one or more uncertain variables whose probability distributions and correlations are given. The validity of probability calculus is undisputed.

There are several interpretations of the concept of probability; some people argue that different interpretations are called for in different contexts. For an exposition of various interpretations, see de Finetti (1968) and essays by de Finetti, Koopman, and Savage in Kyberg and Smokler (1964).

We use an interpretation of the concept of probability that is fundamental to the whole decision analysis philosophy. In this view, probability represents the encoding of knowledge about an uncertain quantity or event. The probability distribution represents a summary of an individual's expert judgment based on all of his experience and informed opinion. This interpretation is certainly valid in situations where one must decide among available alternatives based on

the information at hand; it is compatible with most other interpretations where they apply.

Probability encoding is the process of extracting and quantifying individual judgment about a well defined quantity (which though well defined is uncertain). The purpose of probability encoding is to provide a precise language for communication about uncertainty.

Since various people are likely to have different information, experience, or judgment, two persons can make different probability assignments to the same uncertain quantity. Furthermore, a person is likely to revise a probability assignment if he receives new and relevant information. Indeed, the concept "information" may be defined as anything that causes a person to change a probability assignment; whether a message contains information depends upon what you already know. If it does not alter your probability assignment, you may conclude that the message contains no information for you, i.e., you knew it already, or that the information was irrelevant for the distribution under question.

A probability distribution must meet these three technical conditions:

(1) The probability of any event is a number between zero and one (inclusively).
(2) The probability of an event that is certain to occur is one.
(3) If there are two events that cannot both occur, then the probability that one or the other will occur is equal to the sum of the probabilities for each of the two events.

No other constraints restrict the probability assignment.

A probability distribution can be represented in two different ways. The cumulative probability distribution (cumulative distribution, for short) is one way; an example is shown in Figure 1(a). The possible outcomes for the uncertain quantity are expressed on a value scale (the horizontal axis); the uncertainty in the outcomes is expressed on a probability scale (the vertical axis). Each number on the value scale has a corresponding probability level, which is interpreted as the probability that the resolved value of the uncertain quantity will be less than or equal to that number. (For example, in Figure 1(a), the probability that the value will not exceed 35 is assigned as 0.65.)

An alternative way of describing the distribution is through the density function. The density function has the property that the area under the curve between any two points is equal to the probability that the resolved value will fall between those points. (The total area under the curve is therefore one.) The height of the density function at a particular point thus indicates the relative probability that the resolved value will fall in the vicinity of that point. The density function corresponding to the cumulative distribution of Figure 1(a) is shown in Figure 1(b).

The cumulative distribution and the density function are two equivalent ways of describing a probability distribution. The area under the density function between any two points is always equal to the difference in values of the cumulative distribution for the same points. The steeper the cumulative distribution is, the higher the density function will be at that point on the value scale.

Some Encoding Principles

The following list of principles should be used in defining and structuring any variables whose uncertainty is to be encoded. Violating these principles invariably leads to problems in the

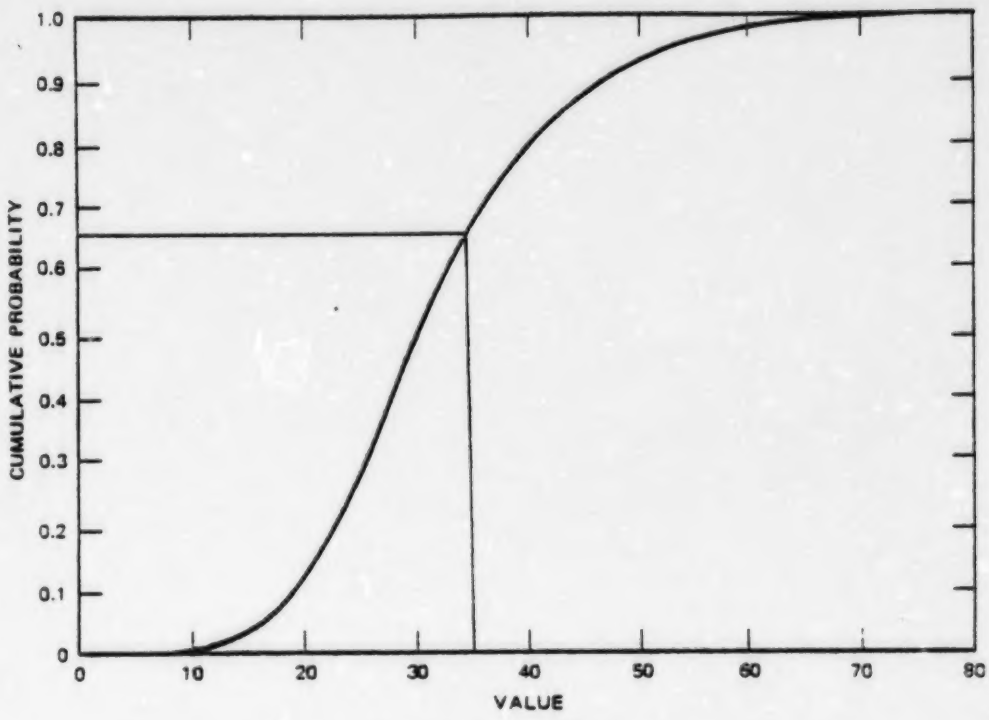

(a) CUMULATIVE DISTRIBUTION

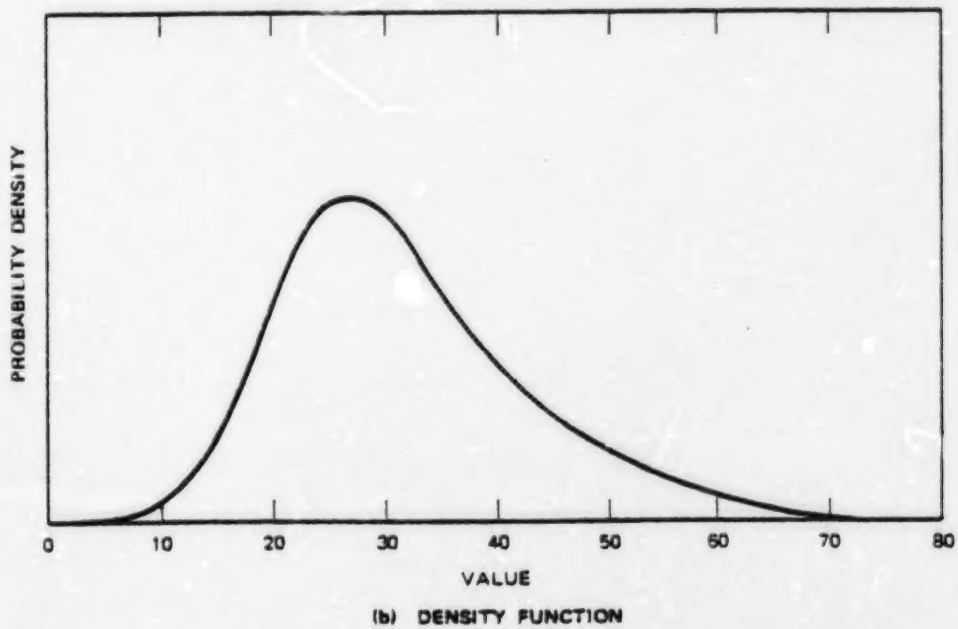

(b) DENSITY FUNCTION

FIGURE 1 TWO REPRESENTATIONS OF A PROBABILITY DISTRIBUTION

probability encoding. The list serves as a checklist befor actual encoding takes place. These principles are:

- The uncertain quantity should be of potential importance to the decision.

- The quantity should be defined for the subject as an unambiguous state variable beyond the subject's control. If the subject believes the outcome of the quantity can be affected to some extent by his decision, then the problem needs restructuring to eliminate this effect.

- The level of detail required in the encoding process depends on the importance of the quantity and should be determined by sensitivity analysis before the interview. It may sometimes be sufficient to elicit only a few points on the distribution.

- The quantity should be well structured. The subject may think of the quantity as conditional on other quantities; accordingly, conditionalities should be considered consciously and brought into the structure because our minds deal ineffectively in combining uncertain quantities. Mental acrobatics should be minimized.

- The quantity should be clearly defined. A good test for clarity is to ask whether a clairvoyant could reveal the value of the quantity by specifying a single number without requesting clarification. To cite an example, it is not meaningful to ask for the "price of wheat in 1974," because the clairvoyant would need to know

the quantity, kind of wheat, at what date, at which Exchange, and whether it is the buying or selling price. However, "the closing price of wheat on June 30, 1975 at the Chicago Commodity Exchange" would be a well defined quantity.

- The quantity should be described by the analyst on a scale that is meaningful to the subject. For example, in the oil industry the subject--depending on his occupation--may think in terms of gallons, barrels, or tank cars. The wrong choice of scale may cause the subject to spend more effort on fitting his answers to the scale than on evaluating his uncertainty. It is important, therefore, to choose a unit with which the subject is comfortable; after the encoding, the scale can be changed to fit the analysis. As a rule, if there is no obvious scale, the subject should choose the scale. Furthermore, choosing the scale for the subject may provide him with information (e.g., about the magnitude of the quantity) that may lead to a bias. It may sometimes be easier to express an uncertain quantity as a fraction of another quantity or as the excess over another quantity rather than in absolute terms.

Balancing Modeling with Encoding

There is always a choice between encoding the uncertainty in an important variable and modeling the problem further. At one extreme, it is conceivable that the value of each alternative could be encoded directly, thus bypassing a need for examination of underlying variables. Generally, a distribution over benefits provides more confidence if a model is constructed that relates benefits to other variables. The modeling effort tends to be most effective and most

economical if it starts with an aggregate model that is successively refined. The model should be refined only while the cost of each addition provides at least comparable improvement in information. This test depends on how the information bears on the decision at hand.

The choice between additional modeling and encoding may need to be reconsidered during the encoding process, because the subject may reveal biases during the interview that could be treated by further structuring of the problem.

Joint Distributions

Most decision problems will have more than one state variable for which the uncertainty should be encoded. (A state variable is not under the control of the decision maker.) Part of the modeling effort is to determine the dependencies among the variables and perhaps to redefine the variables to make new variables that are independent.

Two variables are said to be independent if the probability distribution for one of the variables does not depend on the revealed value of the other variable. For example, the level of ozone that will cause reduced resistance to bacterial infections in sensitive subjects would be judgmentally independent of the level of ozone that would cause a 10% reduction in pulmonary function only if certain knowledge of the one level would add nothing to current knowledge about the other, which seems to be the case. On the other hand, the level of ozone that will reduce athletic performance and the level that will reduce pulmonary function might not be considered independent; indeed, one mechanism for reduction of athletic performance by ozone may be via reduced pulmonary function. If the level that reduced pulmonary function were found to be much lower than current estimates, this could affect experts' judgments on the level causing reductions in athletic performance.

If the variables are modeled as independent, their uncertainties can be encoded separately. Otherwise, the uncertainty in each variable has to be encoded separately for each possible outcome of the variables on which it depends. Dependence increases the labor of the analysis. A discrete example demonstrates this; if each of three variables is modeled to have only three outcomes, thirteen (1 + 3 + 3 x 3) distributions have to be encoded in the case of dependence as compared with only three distributions in the case of independence. The effect of dependence increases rapidly when the number of possible outcomes increases. It is obvious that one should try to model in such a manner that variables are independent. The following paragraph presents an example of how such modeling can be achieved.

A set of variables may be considered dependent because each variable depends in some way on the same underlying variable. However, they may be judgmentally independent for each value of the underlying variable. For example, consider the demands, X1 and X2, for two similar products in a product line in a given period. X1 and X2 are dependent because the two products are competing in the same market. At the same time, it is possible that the first product's share of the market, Y1, is judgmentally independent of the total market, Y2, for the two products. We can then encode the uncertainties for the two new quantities Y1 and Y2 and use the following obvious relationships to get back to the original demands:

$$X1 = Y1 \times Y2 \quad \text{and} \quad X2 = (1 - Y1)Y2 \ .$$

The Encoding Session

It must be emphasized that the encoding of a probability distribution is an art as well as a science. There are indeed certain general principles to keep in mind; but each expert and each session is different.

It is important to have the encoding done in person by a trained analyst aware of the possible difficulties and biases that may arise. The roles of the analyst who elicits the data and the expert whose judgment is being encoded should be separated. The analyst should be aware of the possible types of bias and the difficulties that may be encountered. He should know the laws of probability theory and the relationship between the two equivalent representations of probability. The expert should be able to focus his efforts on making simple judgments regarding the uncertain quantity.

Psychological research has uncovered many possible biases that may cause the judgment expressed by an expert to differ from his true underlying judgment (see Tversky and Kahneman, 1974, and Spetzler and Stael von Holstein, 1975).

It is important to note that these biases are present whenever expert judgment is called on, whether in a formal probability encoding session or in another context. A formal encoding session as described in Spetzler and Stael von Holstein (1975) avoids or corrects for such biases as much as possible. The analyst interviewer, being aware of the common biases, has techniques available to uncover and correct for sources of bias. A simple written questionnaire, on the other hand, may indicate results far from the underlying judgment of the expert who fills it out.

We have found that the actual encoding of judgment takes a relatively small (10-30%) part of the time of a typical session.

The process as practiced by SRI is divided into five phases:

- Motivating--Rapport with the subject is established and possible motivational biases are explored.

- Structuring--The structure of the uncertain quantity is defined.

- Conditioning--The subject is conditioned to think fundamentally about his judgment and to avoid cognitive biases.

- Encoding--This is the actual quantification of judgment in probabilistic terms.

- Verifying--The responses obtained in the encoding are checked for consistency.

Motivating

This phase has two purposes. The first is to introduce the subject to the encoding task. This may entail an explanation of the importance and purpose of the probability encoding and a discussion of the difference between deterministic (single number) and probabilistic (probability distribution) judgments.

The second purpose is to explore whether any motivational biases might operate. The interviewer and the subject should have an open discussion on what rewards might be associated with the probability assignment, as well as on possible misuses of the same information. The subject may be aware of misuses of single-number predictions, e.g., that they often are interpreted as "commitments." It should be pointed out that no commitment is inherent in a probability distribution. In fact, the distribution represents the complete judgment of the subject. During the discussion, the subject is likely to reveal some reactions indicating biases that might be expected later in the encoding process. For example, the subject may display some caution against over-estimation and may, therefore, be somewhat biased toward under-estimation. The planned use of the encoded information should be discussed. It should be stressed that the expert should not include a 'margin of error' in his judgments; on the other hand, he should take account of the uncertainties in current

experimental information. (The relevance of some current information to the quantity being encoded may also be a matter of judgment.) What is wanted is the expert's best judgment regarding the quantity in question, including the uncertainty regarding the quantity.

Structuring

The next step in the encoding process is to define and structure clearly the uncertain quantity. This quantity is assumed to be important to the decision. It should be defined as an unambiguous state variable. The definition should pass the clairvoyant test, i.e., a clairvoyant should be able to specify the outcome without asking additional questions for clarification. The structure may have to be expanded so that the subject does not have to model the problem further before making each judgment. It is also important to choose a scale that is meaningful to the subject.

The subject should be required to think the problem through carefully before the actual encoding session begins. He should decide what background information might be relevant (or irrelevant) to the problem. Otherwise, only the readily available information will be used initially, and new information may surface later in the session and invalidate all prior answers. Even if it does not, however, the resulting distribution may be biased with respect to the subject's underlying judgment.

Conditioning

The aim of this phase is to head off biases that otherwise might surface during the encoding and to condition the subject to think fundamentally about his judgment. Basically, the phase should be directed toward finding out how the subject goes about making his probability assignments. This will reveal what information seems to be most available, what (if any) anchors are being used, what assumptions are made, and so on. The interviewer should thus watch

out for (as well as make use of) the modes of judgment discussed in Tversky and Kahneman (1974) and Spetzler and Stael von Holstein (1975). This could be done in the form of a checklist. The development of an efficient general checklist is subject to further research, but the following suggestions have been found useful in many applications.

The subject can be asked to specify the most important bases for his judgment. With air pollution standards, these may be 'consensus' values as expressed in conference reports or in previous standards. Such values could then be expected to act as anchors and often lead to a central bias.

Asking the subject what he is taking into account will show what information becomes most easily available. The interviewer can also make use of availability if he suspects the subject to have a central bias. He can then ask the subject to compose scenarios that would produce extreme outcomes. The subject could be asked where in his experience he had really been surprised when the uncertainty had been resolved for some important quantity, or if any similar cases had surprised him. This might lead to the general question of what could lead to a surprise outcome for the quantity currently being encoded.

Sometimes, the uncertain quantity is an average, such as average productivity or average reliability. The interviewer should then try to determine whether the distribution assigned by the subject really is a distribution for the average or a distribution for an individual unit. (The reason is that people often have difficulty in discriminating between the two situations.) If the latter is the case, it is probably best to use the resulting distribution and restructure the model. Representativeness may come into play in another situation when one is concerned with revising a probability assignment in the light of new information. The best way to handle such a situation is often to ask for the probability distribution of the quantity without the new information and for the probability of

the information conditional on the outcome of the quantity; it is then a matter of applying probability calculus to obtain the distribution for the quantity given the information. This procedure was used in Howard, Matheson, and North (1972).

It is important to specify all assumptions (conditionalities) that may underly the probability distribution, as well as those factors that the subject is supposed to integrate into his judgment. These may have been considered when the quantity was defined and structured in the preceding phase. Otherwise, the subject should be asked if there are any assumptions that he has made because he does not feel responsible for taking account of the corresponding uncertainties. This may sometimes lead to a change in the structure because some conditionalities have been made explicit. The encoding may then be made conditional on different sets of assumptions, and the probabilities that the various assumptions will hold are then encoded separately (from the current subject or from someone else).

When a subject is assigning a probability to the occurrence (or nonoccurrence) of some event (e.g., that a product will be successful in the market or that a given ambient ozone concentration will actually cause a given health effect), he may base his assignment on whether he can generate plausible scenarios leading to the occurrence of the event in question. (A 'scenario' may be a sequence of research results.) Asking him to state the basis for his probability assignment may reveal that the coherence of such scenarios has been an important factor. The interviewer may then want to generate more scenarios that would or would not lead to the occurrence of the event. For example, simply devising a scenario that implies the opposite outcome might change the first probability assignment.

Encoding

The first three phases have defined the quantity for which the uncertainty is to be encoded, the structure underlying the quantity,

and the scale to be used for the quantity. They have brought forward any cognitive biases that may be operating and have hopefully eliminated or greatly reduced the motivational biases. The procedure outlined for this phase of the interview process is suggested as a guideline for the actual quantification of judgment in probabilistic terms. It is primarily based on the use of the probability wheel as the encoding technique. Responses may often indicate a need to return to the tasks in the previous three phases.

The interviewer begins by asking for extreme values for the uncertain quantity. Then he asks for scenarios that might lead to outcomes outside the extremes. Also ask for the probabilities (or odds) of outcomes outside the extremes. This deliberate use of availability is to counteract the central bias that is otherwise likely to occur. Eliciting the scenarios makes them available to the subject, who is then likely to assign higher probabilities to extreme outcomes; this has the effect of increasing the variability in his assigned distribution.

The actual encoding makes use of a "probability wheel." This is a device containing a pie-chart of two colors, such as red and blue. The relative amount of red may be adjusted. The analyst asks the subject for judgments of which of two events is more likely, where one event involves spinning the wheel and the other involves the quantity being encoded. In this way, the subject has only to make a simple judgment of relative likelihood.

It is important to note that the subject is not led by the hand in any way. The analyst is careful not to suggest values that are 'good' or 'bad.' The subject himself provides the scale in which values are expressed, as well as the range of values encoded. In the main encoding process, the subject is not asked directly for probability numbers, but is asked only to tell which of two possible events is more likely.

The interviewer should be aware of attention shifts (for example, the shift of attention between the encoding process and the actual problem), changes in the subject's modeling of the situation, and the appearance of new information. In the middle of the encoding process, the subject may remember pieces of information that may bear on the outcome of the uncertain quantity. Such information may invalidate all points obtained so far in the process. The encoding process is time-consuming, and it may be difficult to maintain a high motivation for the subject. It often helps to keep his interest, however, if he is shown some inconsistencies among his responses.

Each response leads to a point on a cumulative distribution. The importance of the variable for the decision problem at hand determines the number of points to encode. The distribution is not shown to the subject during the encoding process; he is encouraged to give his best judgment at each point, and not to be worried about overall consistency. The consistency of the resulting points is then one measure of the precision and stability of the subject's judgement. After enough points have been encoded, a curve should be fitted to the points. An example is shown in Figure 2.

Verifying

In the last phase, the judgments are tested to see if the subject really believes them. If needed, there may be iteration through some of the stages mentioned above.

A graphical representation of the responses as points on a cumulative distribution and an interpretation of this distribution (perhaps in terms of a density function) provides an important test and feedback. An examination of the distribution itself cannot show whether or not the distribution agrees with the subject's judgment. However, it can show implications of the subject's responses and thereby provide feedback. If some responses are not consistent with the subject's judgment, they will have to be modified.

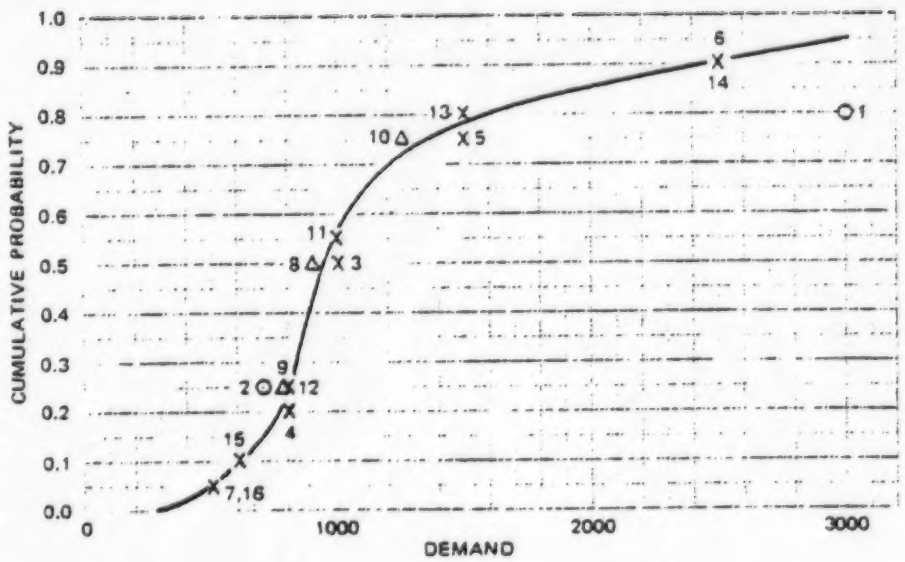

FIGURE 2 EXAMPLE OF A CURVE FITTED TO RESPONSES

The analyst should explore the logical consequences of the encoded distribution with the subject, and confirm that the encoded distribution represents his best judgment. For instance, the general shape of the curve should be noted (is it bimodal? Are there long tails?). Another technique for verification is to construct fair bets, chosen so that each bet is equally attractive if the encoded distribution represents the subject's underlying judgment.

III THE QUESTION OF VALIDATION

A judgmental probability distribution is valid if it is a true representation of the expert's judgment. A judgmental probability distribution depends on the information and judgment of the expert. If new information, for instance, new experimental results, is made available, one would expect an expert's judgment about a quantity, and thus his judgmental probability distribution, to change. This is by no means a weakness of the method. A judgmental probability distribution represents current uncertainty; new information can reduce uncertainty.

It is thus difficult to find an 'objective' validation for a judgmental probability distribution--that is, a validation that can be made in a different way than returning to the expert to ask his opinion. In one-time, nonrepetitive instances, the best way to judge the validity of the procedure is to ask those concerned if they felt the judgmental probability approach helped clarify and quantify their judgments.[5]

Several studies have been made in repetitive situations of the accuracy of particular judgmental probabilities. In general, the judgmental probabilities have been borne out by future events, as the following examples show.

Weather Forecasting

Murphy and Winkler (1977) have studied weather forecasters' assignments of probabilities. They find that weather forecasters do a very good job. (Of course, they have the benefit of feedback to enable them to improve their performance.) For example, Figure 3 compares the forecast probability of precipitation with the observed relative frequency. It is based on all 17,514 forecasts issued by the

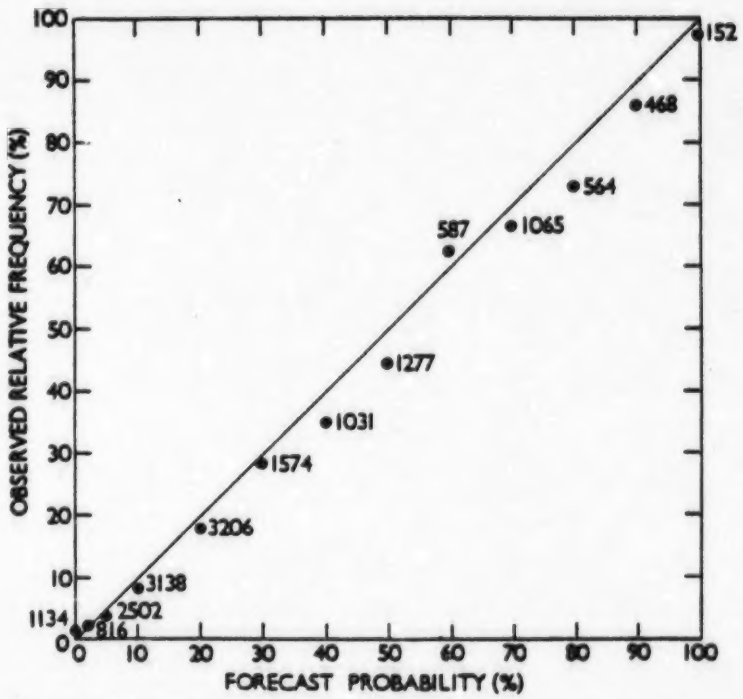

SOURCE: Murphy and Winkler, 1977

FIGURE 3 THE RELIABILITY DIAGRAM FOR ALL OF THE PRECIPITATION PROBABILITY FORECASTS FORMULATED BY NWS FORECASTERS AT CHICAGO DURING THE PERIOD FROM JULY 1972 TO JUNE 1976

National Weather Service in Chicago from July 1972 to June 1976. Th numbers are the number of forecasts in each interval. The diagonal line represents complete reliability, in the sense that the relative frequencies exactly equal the forecast probabilities. As Murphy and Winkler say, "Examination of this figure reveals that these forecasts were very reliable; that is, the observed relative frequencies corresponded very closely to the forecast probabilities."

Future Interest Rates

Since 1971, Morgan Guaranty Trust has used subjective probability distributions in the form of 'histograms' (a form of the probability density function) to treat uncertainty in future interest rates. A track record of this activity was reported by Kabus (1976), senior operations research officer. This track record was quite good. Kabus compared monthly predicted interest rates over 1 year with the actual values attained by the rates and concluded

> The forecasted and actual values of the 90-day CD rate for the horizon dates (were) fairly close together. Further, the 90-day CD rate attained no surprising level (outside the 75% confidence interval) on any of the horizon dates. Finally, the direction (of interest rate movement) was always forecast correctly.

The probability distributions were used to help decide between purchases of 90-day or 180-day certificates of deposit (CDs). Kabus was quite happy with the outcomes of the decisions: "Exhibit VI clearly indicates through the proximity of the cost of money from the 'money curve recommendation' to that of the 'best' situation (4.71% versus 4.70%) that the expectations implied by the 90-day CD histograms were indeed quite accurate."

His report concluded with the following note.

> Histograms have been a part of Morgan Guaranty's asset/liability decision-making process for the past five years. It appears likely that they will continue to serve as the means by which top executives are made aware of how much uncertainty underlies the expectations of future interest rates that are presented to them. One major reason for the success of the histogram at Morgan Guaranty has been the complete understanding that the only "incorrect" histogram is one which does not correctly reflect its creator's expectations. While there is no guarantee that histograms will give an accurate picture of the future, they will give an accurate picture of an individual's view of the future. Thus, if the histogram of some future event changes with time, it is only because people change their opinions with time, not that the previous histogram was incorrect.

R&D Outcomes

Sandoz Pharmaceuticals has used the quantification of expert judgment as a part of its R&D planning since 1970. Figure 4, taken from Balthasar, Boschi, and Menke (1978), compares 3 years of expert judgments on R&D technical success with the actual results 1 to 5 years later. (It is similar to Figure 3 in concept.) Judgmental probabilities of technical success on pharmaceutical R&D projects were assigned by an R&D manager, who consulted individual experts. The horizontal axis plots the assigned probabilities; the vertical axis is the actual observed frequency of success for those products predicted to have a given success probability. Complete accuracy of predictions would imply that the results lie along the diagonal line. The actual results, shown by the black dots, are close to the horizontal line. In all but one case, they lie within the vertical error bars (one standard

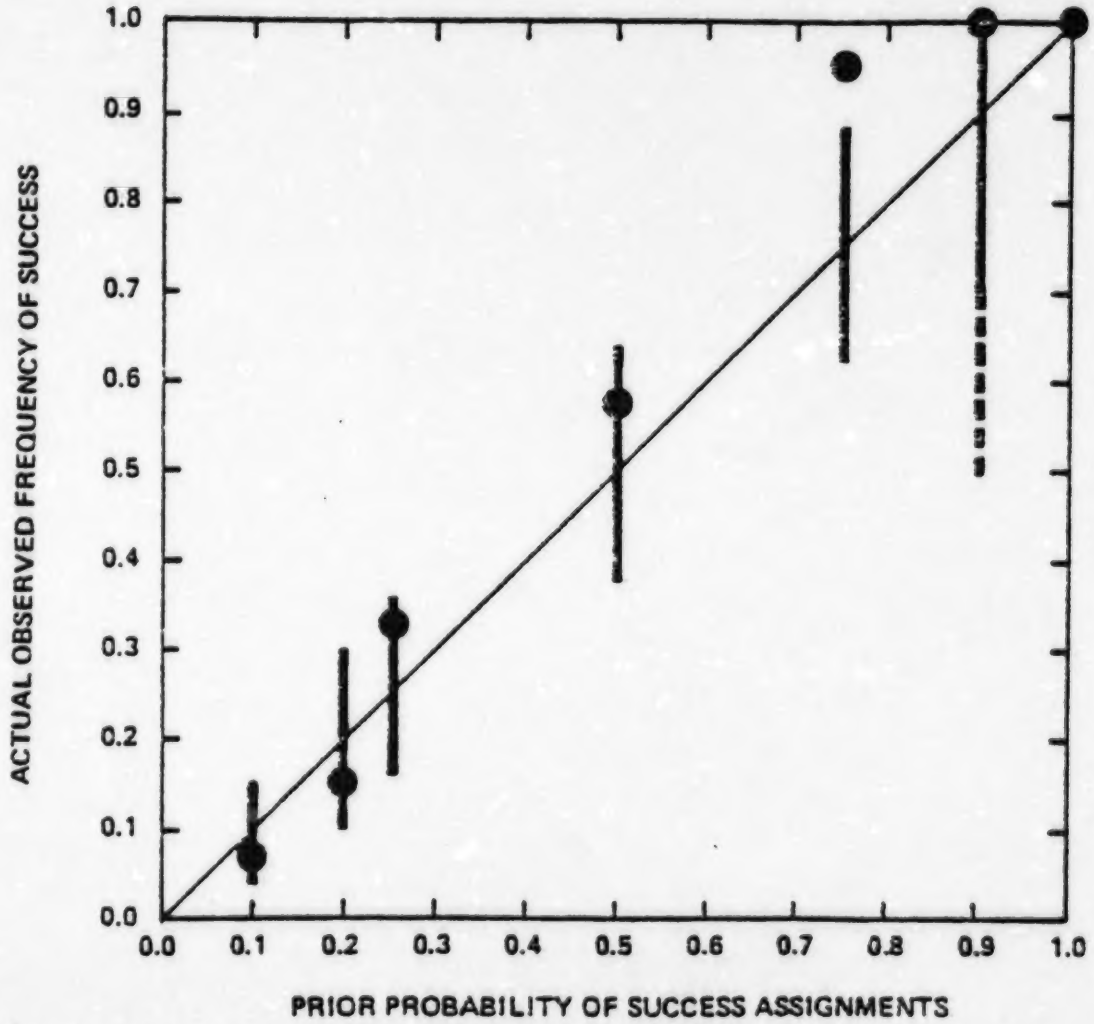

*SOURCE: Balthasar, Boschi, and Menke, 1978

FIGURE 4 RELIABILITY OF EXPERTS' JUDGEMENT (1970-72)*

deviation). The data are consistent with perfect calibration, that is, perfect validation of the probability estimates on a statistical basis.

This project has proved so useful to Sandoz Pharmaceuticals that forecasts are now being made by an expert panel consisting of about a dozen R&D department heads and managers, as well as some clinical specialists.

Review Articles and Bibliographies

Slovic, Fischoff, and Lichtenstein (1976) and Hogarth (1975) have written review articles with extensive bibliographies. See also Wendt and Vlek (1975) and Jungerman and Zeeuw (1977).

FOOTNOTES

[1] The term, 'subjective probability' has often been used. This term caries the unwarranted connotation that there are two kinds of probability, subjective and objective. According to our view, there is only one kind of uncertainty. Uncertainty is a property of a person's state of knowledge; in all cases, a person must use his best judgment. This includes the use of the best, most appropriate 'objective' data where available.

Professor Myron Tribus (1978) of MIT, former Assistant Secretary of Commerce, also feels the use of the term 'subjective' in this context unnecessarily opens the method to unwarranted criticism. He prefers a term such as 'encoding people's knowledge' to 'encoding a subjective probability distribution.'

For an exposition of this interpretation of probability, and a comparison with other views, see de Finetti (1968).

[2] Concerning the probability of the stock depletion of offshore spotted porpoises, Eastern Spinner porpoises, and Whitebelly Spinner porpoises, Dr. Robert M. White, Administrator, National Oceanic and Atmospheric Administration, said, during hearings before the Subcommittee on Fisheries and Wildlife Conservation and the Environment of the Committee on Merchant Marine and Fisheries (see U.S. Government Printing Office, 1976):

> A porpoise stock may be determined to be depleted when its population falls below the lower limit of the 'optimum sustainable population' (OSP). Theoretically, the determination of depletion is a straightforward matter. The stock size is ascertained and compared with the OSP level. Unfortunately, both quantities are subject to considerable

uncertainty in their possible ranges and cannot easily be compared by inspection. This uncertainty can be represented, however, by considering each quantity as a random variable. It is then possible to determine the probability that stock size is below the OSP level.

A probabilistic representation was undertaken of the state of knowledge expressed in the recent STock Assessment Report about optimum sustainable population and the three porpoise stocks. If OSP is in the range of 60-70%, then the probability of the Offshore Spotted stock being depleted is 50%, and the probability of the Eastern Spinner population being below OSP is 73.6%. For the Whitebelly Spinner, if the OSP range is 60-70%, the probability is 16.3% that the stock is depleted.

3 Philip Handler, President of the National Academy of Sciences, said in his letter of transmittal of Air Quality and Stationary Source Emission Control (see North and Merkhofer, 1975).

Perhaps the principal contribution of the report is to reveal what is and is not known, and in this way guide both the decisions which must be taken in the face of uncertainty and the future research program necessary to reduce that uncertainty.

4 People perceive and assess the magnitudes associated with various quantities in a manner that is analogous to the manner in which they perceive and assess distances. When they do not have access to detailed information or direct measurements, they often have access to cues of limited validity and reliability that serve as bases for an estimation. In making these intuitive estimates, people follow

unconscious procedures called modes of judgment. More often than not, these procedures lead to reasonable answers. To survive together, drivers and pedestrians have learned to estimate distances and speeds with considerable accuracy. Likewise, the security analyst has learned to estimate with accuracy the likelihood of success of business enterprises. Although the intuitive rules often yeild reasonable estimates, they can also be subject to systematic biases that are severe and costly.

In estimating distances, for example, people tend to overestimate the distance of a remote object when visibility is poor, and to underestimate the distance when the sky is clear. This bias arises because the haze through which a feature of the landscape is seen is a cue to its distance. This cue has some validity because more distant objects are usually seen through more haze. Under specified circumstances, such as on exceptionally foggy or clear days, reliance on this cue leads to predictable errors of judgment. Three features of this example are worth noting. (1) People are not generally aware of the cues on which their judgments are based. Few people consciously realize that they use haze to judge distances, although research shows that virtually everybody does. (2) It is difficult to control the cues we use; the object seen through haze looks more distant, even when we know why. (3) People can be made aware of the bias and can make a conscious attempt to control its effects--the captain of a ship does this when navigating in a mist.

An analogous problem exists in the assessment of uncertain quantities. Here too, one relies on certain modes of judgment that may introduce systematic biases. Here too, modifying impressions and intuitions is exceedingly difficult, but it is possible to learn to recognize the conditions under which such impressions are likely to lead us astray. See Tversky and Kahneman (1974).

5 Business users have often expressed satisfaction with decision analysis and the subjective probability encoding that is a part of

it. We quote from the closing of a speech by William C. Roher, President, Gulf Oil Chemicals Association, before the 33rd annual Honor Award meeting, Commercial Development Association (March 9, 1977), which was devoted to a discussion of a decision analysis of a new olefins plant performed with SRI International help: (Quoted by permission of Mr. Roher)

> I shall conclude with one more principle, which I am happy to borrow from Sir Francis Bacon, and I quote:
>
> > 'If a man will begin with certainties he will end with doubts, but if he will be content to begin with doubts he shall end in certainties.'
>
> We in Gulf Oil Chemicals saw during our olefins analysis that a business we knew very well had its share of uncertainties and doubt. We also saw that in decision analysis we had a tool for making those uncertainties explicit. With that tool, we could communicate to all concerned--in easily understood terms of profits--the magnitude of both the risk and the opportunity associated with those uncertainties.
>
> We had begun in doubt. With the help of decision analysis, we ended with certainty--the certainty of having made a good decision. Thus, if you feel as I do that good decisions are more often than not followed by good outcomes, you should not hesitate to use decision analysis where the outcome is important.

REFERENCES

Balthasar, H. U., R. A. A. Boschi, and M. M. Menke, "Calling the Long Shots in R&D," Harvard Business Review, Vol. 56, No. 3, pp. 151-160 (May-June 1978).

Blue T.A., et al., New, Innovative Pesticides: An Evaluation of Incentives and Disincentives for Commercial Development by Industry, SRI International, Menlo Park, California (May 1977).

Buckbinder, B., S. G. Helzer and F. L. Offensend, "Preliminary Report on Evaluating Alternatives for Reducing Upholstered Furniture Fire Losses," NBSIR 77-1381, National Bureau of Standards, U.S. Department of Commerce (November 1977).

Cazalet, E. G., B. R. Judd, A. C. Miller, III, and E. S. Rasmussen, Decision Analysis of California Electrical Capacity Expansion, report prepared for the Energy Assessments Division of the California Energy Resources Conservation and Development Commission, SRI International, Menlo Park, California (February 1977).

Cazalet, E. G., et al., "Recommendations for a Synthetic Fuels Commercialization Program," in "Cost/Benefit Analysis of Alternative Production Levels," Report #041-001-00111-3, submitted by the Synfuels Interagency Task Force to the President's Energy Resources Council, U.S. Government Printing Office, (November 1975).

de Finetti, B., "Probability: Interpretations," in International Encyclopedia of the Social Sciences, (ed.) D. L. Sills, Vol. 12, pp. 496-505 (Macmillan and Free Press, New York, 1968).

EPA Staff, A Method for Assessing the Health Risks Associated with Alternative Air Quality Standards for Photochemical Oxidants, Strategies and Air Standards Division, Office of Air Quality Planning and Standards, Office of Air and Waste Management, U.S. Environmental

Protection Agency, Research Triangle Park, North Carolina (Jan. 6, 1978). A final version of this paper is scheduled to appear in May, 1978.

EPA, "Oxidant/Ozone Health Effects Air Quality Criteria Document," Environmental Protection Agency (1978).

Harrison, J. M., D. W. North, and C. A. S. Stael von Holstein, Decision Analysis of Wildland Fire Protection: A Pilot Study, report prepared for the Forest Service, U.S.D.A., SRI International, Menlo Park, California (February 1973).

Hogarth, R. M., "Cognitive Processes and the Assessment of Subjective Probability Distributions," J. Am. Stat. Assoc., Vol. 70, No. 350, pp. 271-294 (June 1975).

Howard, R. A., "The Foundations of Decision Analysis," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, pp. 200-210 (September 1968).

Howard, R. A., "Decision Analysis: Applied Decision Theory, "Proc. Fourth Inter. Conf. on Oper. Research, pp. 55-71 (Wiley-Interscience 1966).

Howard, R. A., J. E. Matheson, and D. W. North, "The Decision to Seed Hurricanes," Science 176, pp. 1101-1202 (1972).

Howard, R. A., J. E. Matheson, and D. W. North, Decision Analysis for Environmental Protection Decisions, prepared for the National Academy of Sciences, Washington, D.C. (June 1977).

Judd, B. R., D. W. North, and J. P. Pezier, "Assessment of the Probability of Contaminating Mars," SRI International, Menlo Park, California (February 1974).

Jungermann, H., and G. de Zeeuw (eds), Decision Making and Change in Human Affairs, Reidel, Dordrecht, The Netherlands (1977).

Kabus, I., "You Can Bank on Uncertainty," Harvard Business Review, pp. 95-105 (May-June 1976).

Kyberg, H.E., and H.E. Smokler (eds.), Studies in Subjective Probability (Wiley, New York 1964).

Lawrence Livermore Laboratories, UCRL-52303 (1977).

Matheson, J. E. and W. J. Roths, Decision Analysis of Space Projects: Voyager Mars, paper presented at the national symposium, "Saturn/Apollo and Beyond," American Astronautical Society (June 11, 1967).

Merkhofer, M. W., and R. J. Korsan, Florida Utility Pollution Control Options and Economic Analysis, Volume 2: Cost-Benefit Analysis of Alternative Florida Sulfur Oxide Emissions Control Policies, SRI International, Menlo Park, California (January 1978).

Merkhofer, M. W. and D. J. Quinn, Methodology of Back-Contamination Risk Assessment for a Mars Sample Return Mission, SRI International, Menlo Park, California (April 1977).

Murphy, A. H. and R. L. Winkler, "Reliability of Subjective Probability Forecasts of Precipitation and Temperature," J. of the Royal Statistical Society 25, pp. 41-47 (1977).

National Academy of Sciences, Decision Making for Regulating Chemicals in the Environment, National Academy of Sciences, Washington, D.C. (1975).

National Research Council, Decision Making in the Environmental Protection Agency, National Research Council, Washington, D.C. (1977).

North, D. W. and M. W. Merkhofer, "Analysis of Alternative Emissions Control Strategies," in Air Quality and Stationary Source Emission Control, a report by the Commission on Natural Resources, National Academy of Sciences, National Academy of Engineering, and National Research Council, prepared for the Committee on Public Works, U.S. Senate, March 1975, Serial No. 94-4, U.S. Government Printing Office, Washington, D.C. (1975).

North, D. W. "A Tutorial Introduction to Decision Theory," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-4, No. 3, pp. 200-210 (September 1968).

Offensend, F. L., and M. W. Merkhofer, Elements of a Planning Methodology for Use by the National Caries Program, SRI International, Menlo Park, California (August 1976).

Palmes, E.D., "Dose-Response Relations," in Air Quality and Automobile Emission Control, pp. 479-481, a report by the Coordinating Committee in Air Quality Studies, National Academy of Sciences, prepared for the Committee on Public Works, United States Senate, U.S. Government Printing Office, Washington, D.C. (1974).

Raiffa, H., Decision Analysis: Introductory Lectures on Choices Under Uncertainty (Addison-Wesley, 1968).

Slovic, P., B. Fischoff, and S. Lichtenstein, Behavioral Decision Theory, Technical Report DDI-7 by the Oregon Research Institute for the Defense Advanced Research Projects Agency, September 1976. Address inquiries to Dr. Martin A. Talcott, Director, Engineering Psychology Programs, Office of Naval Research, 800 N. Quincy St., Arlington, VA, 22217.

Spetzler, C. S. and C.A.S. Stael von Holstein, "Probability Encoding in Decision Analysis," Management Science, Vol. 22, No. 3 (November 1975).

SRI International, "Readings in Decision Analysis," second edition, SRI International, Menlo Park, California (1977).

Synfuels Interagency Task Force to the President's Energy Resources Council, "Synthetic Fuels Commercialization Program," Vol. II: Cost/Benefit Analysis of Alternate Production Levels (November 1975).

Tribus, M., private communication (February, 1978).

Tribus, M., Rational Descriptions, Decisions, and Designs (Pergamon Press, 1969).

Tversky, A. and D. Kahneman, "Judgment Under Uncertainty: Heuristics and Biases," Science 185, pp. 1124-1131 (1974).

U.S. Government Printing Office, Whales and Porpoise Hearings, printed for the House Committee on Merchant Marine and Fisheries, Serial No. 94-95 (September 1976).

Wendt, D. and C. A. J. Vlek, (eds), Utility, Probability, and Human Decision Making, Reidel, Dordrecht, The Netherlands (1975).